Scalable Supernova Simulation with CHIMERA

NATIONAL CENTER
FOR COMPUTATIONAL SCIENCES

Bronson Messer
Scientific Computing Group
NCCS
&
Theoretical Astrophysics Group
Oak Ridge National Laboratory

Co-authors and collaborators

- Steve Bruenn (Florida Atlantic University)
- Raph Hix (ORNL)
- John Blondin (North Carolina State University)
- Tony Mezzacappa (ORNL)

- Suzanne Parete-Koon (UTK)
- Ted Lee (UTK)

Outline

- Some supernova science
- Why is it a petascale (and beyond) problem?
- CHIMERA* architecture
- Some testing results

* Conservative Hydrodynamics Including Multi-Energy RAdiation

Supernova modeling marked the genesis of computational astrophysics

THE HYDRODYNAMIC BEHAVIOR OF SUPERNOVAE EXPLOSIONS*

STIRLING A. COLGATE AND RICHARD H. WHITE
Lawrence Radiation Laboratory, University of California, Livermore, California

Received June 29, 1965

ABSTRACT

We regard the release of gravitational energy attending a dynamic change in configuration to be the primary energy source in supernovae explosions. Although we were initially inspired by and agree in detail with the mechanism for initiating gravitational instability proposed by Burbidge, Burbidge, Fowler, and Hoyle, we find that the dynamical implosion is so violent that an energy many times greater than the available thermonuclear energy is released from the star's core and transferred to the star's mantle in a supernova explosion. The energy released corresponds to the change in gravitational potential of the unstable imploding core; the transfer of energy takes place by the emission and deposition of neutrinos.

"The reason this paper is cited so many times is **because it started the new endeavor of hydrodynamic stellar modeling.** It is ironic that this work started because of an argument with Soviet scientists during the negotiations for the Cessation of Nuclear Weapons Tests in Geneva in 1959. It was claimed by me that the radiation emissions from a supernova might trigger the then proposed detection net for high altitude nuclear explosions that the Soviets were proposing. This objection of a possible false triggering of the system was brushed aside by the Soviet Ambassador Tsarpkin because, 'Who knows what a supernova would look like?'"

- S. Colgate *The Scientist* 12/1/1980

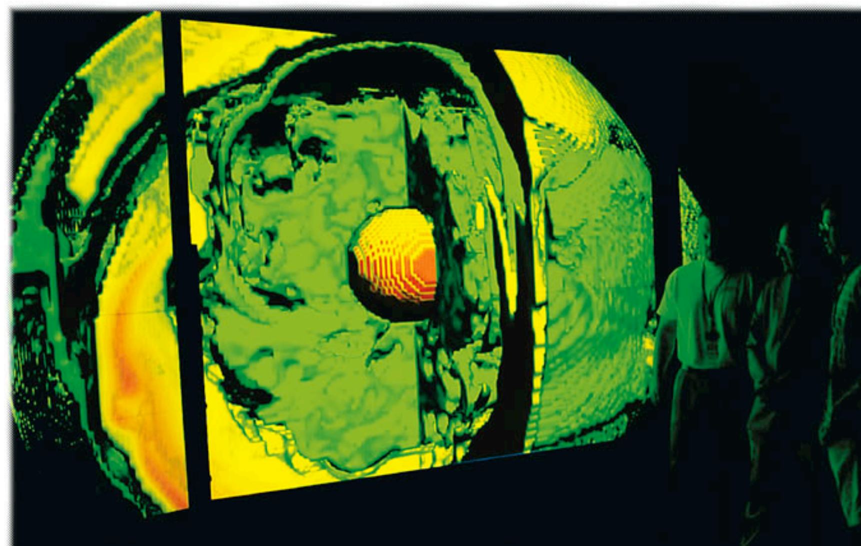

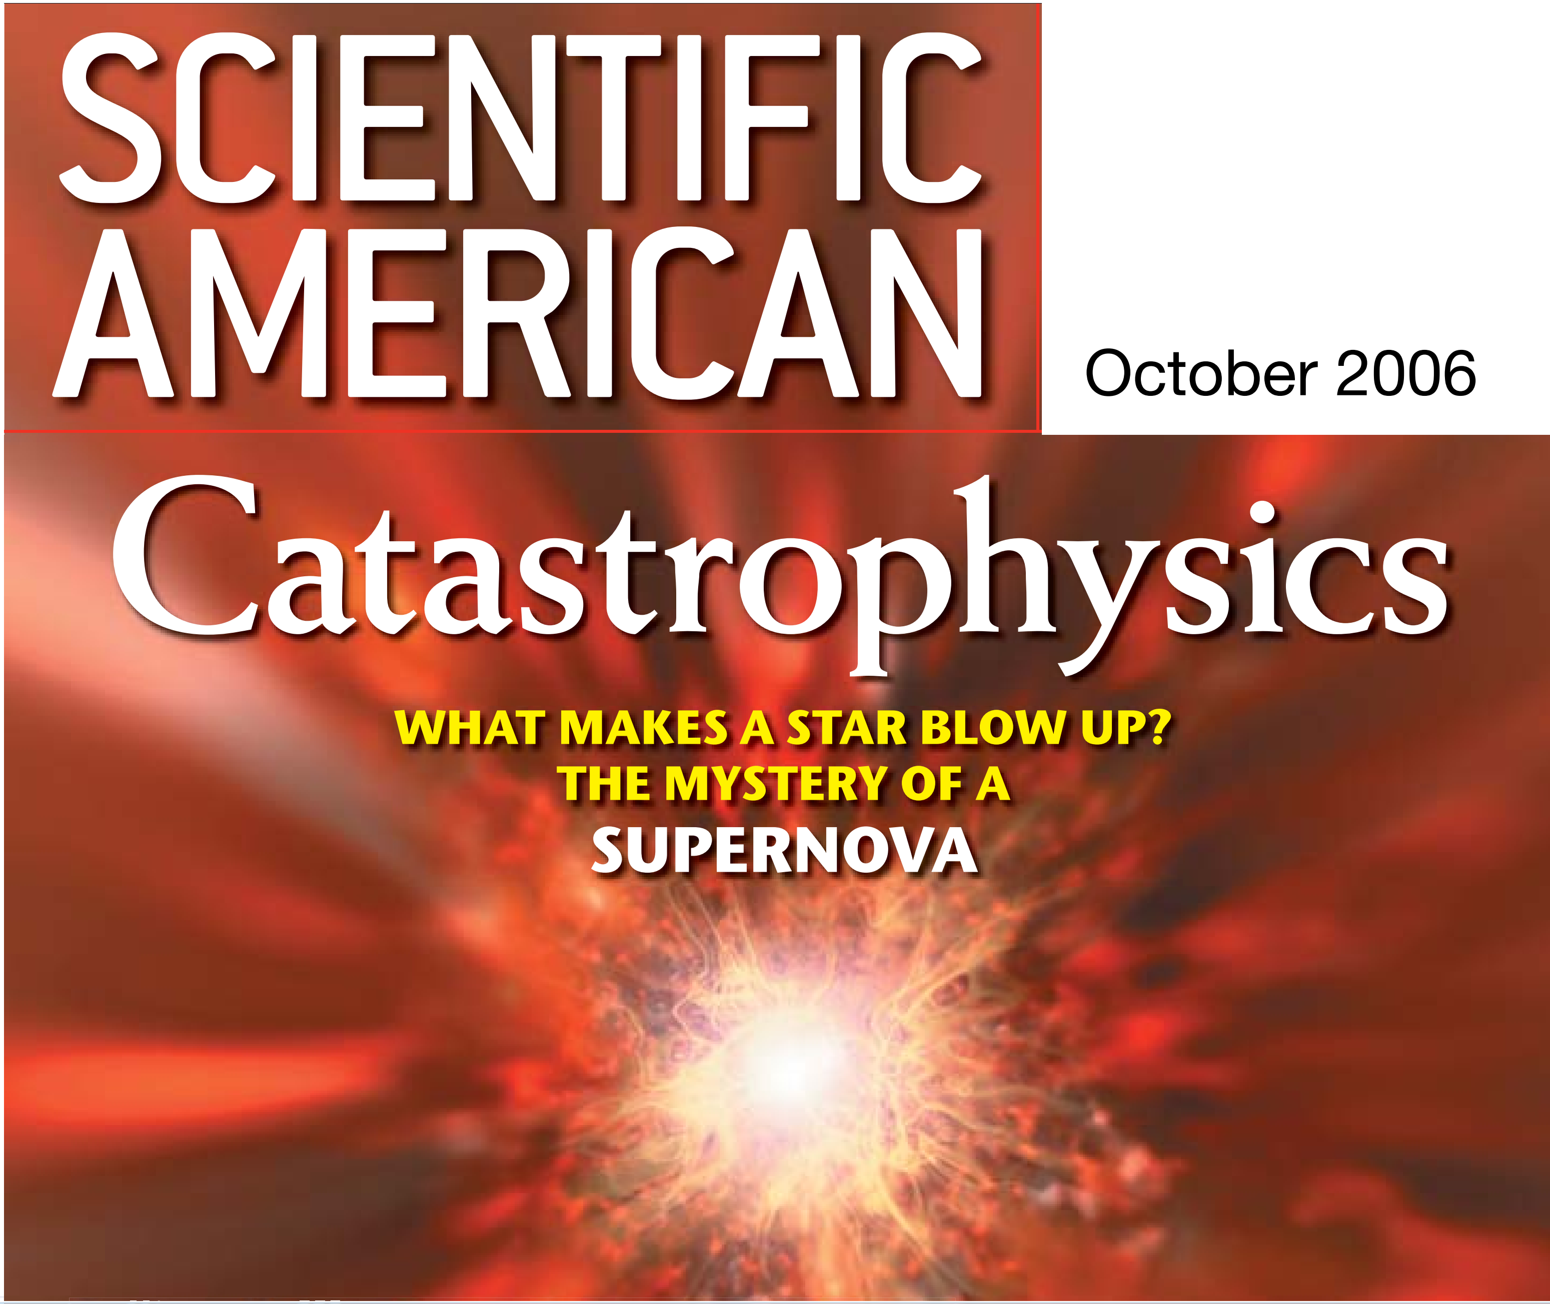

October 2006

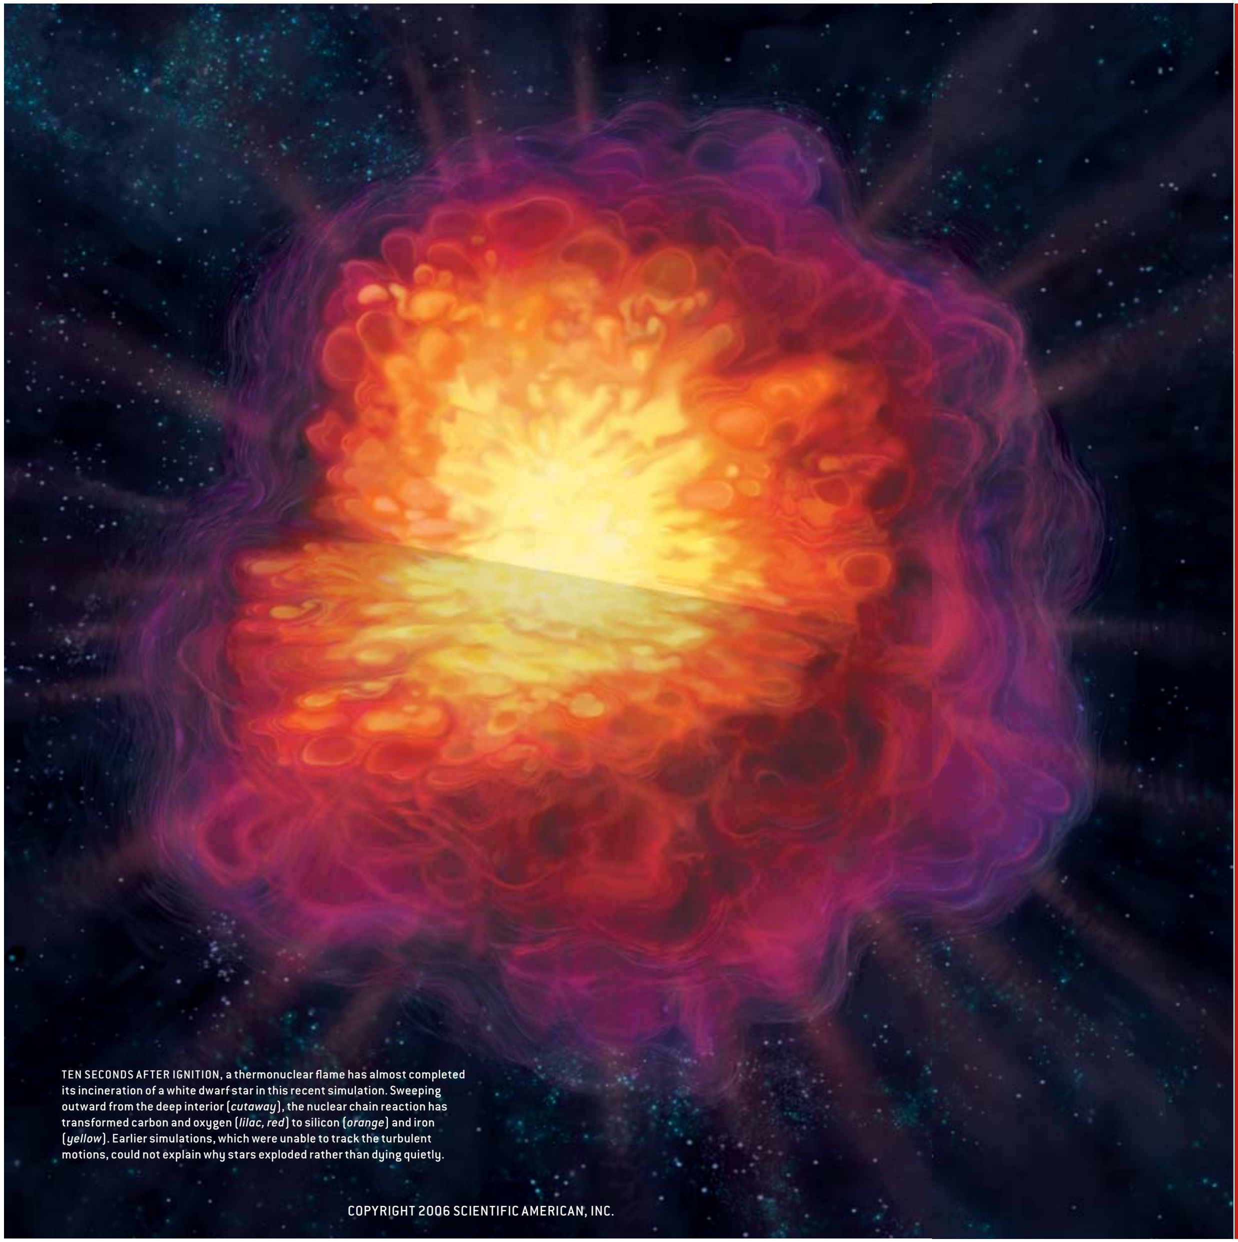

TEN SECONDS AFTER IGNITION, a thermonuclear flame has almost completed its incineration of a white dwarf star in this recent simulation. Sweeping outward from the deep interior (*cutaway*), the nuclear chain reaction has transformed carbon and oxygen (*lilac, red*) to silicon (*orange*) and iron (*yellow*). Earlier simulations, which were unable to track the turbulent motions, could not explain why stars exploded rather than dying quietly.

How to Blow Up a Star

By Wolfgang Hillebrandt, Hans-Thomas Janka and Ewald Müller

It is not as easy as you would think. Models of supernovae have failed to reproduce these explosions—until recently

On November 11, 1572, Danish astronomer and nobleman Tycho Brahe saw a new star in the constellation Cassiopeia, blazing as bright as Jupiter. In many ways, it was the birth of modern astronomy—a shining disproof of the belief that the heavens were fixed and unchanging. Such "new stars" have not ceased to surprise. Some 400 years later astronomers realized that they briefly outshine billions of ordinary stars and must therefore be spectacular explosions. In 1934 Fritz Zwicky of the California Institute of Technology coined the name "supernovae" for them. Quite apart from being among the most dramatic events known to science, supernovae play a special role in the universe and in the work of astronomers: seeding space with heavy elements, regulating galaxy formation and evolution, even serving as markers of cosmic expansion.

Zwicky and his colleague Walter Baade speculated that the explosive energy comes from gravity. Their idea was that

TEN SECONDS AFTER IGNITION, a thermonuclear flame has almost completed its incineration of a white dwarf star in this recent simulation. Sweeping outward from the deep interior (*cutaway*), the nuclear chain reaction has transformed carbon and oxygen (*lilac, red*) to silicon (*orange*) and iron (*yellow*). Earlier simulations, which were unable to track the turbulent motions, could not explain why stars exploded rather than dying quietly.

How to Blow Up a Star

By Wolfgang Hillebrandt, Hans-Thomas Janka and Ewald Müller

It is not as easy as you would think.

have failed to reproduce these explosions—until recently

On November 11, 1572, Danish astronomer and nobleman Tycho Brahe saw a new star in the constellation Cassiopeia, blazing as bright as Jupiter. In many ways, it was the birth of modern astronomy—a shining disproof of the belief that the heavens were fixed and unchanging. Such "new stars" have not ceased to surprise. Some 400 years later astronomers realized that they briefly outshine billions of ordinary stars and must therefore be spectacular explosions. In 1934 Fritz Zwicky of the California Institute of Technology coined the name "supernovae" for them. Quite apart from being among the most dramatic events known to science, supernovae play a special role in the universe and in the work of astronomers: seeding space with heavy elements, regulating galaxy formation and evolution, even serving as markers of cosmic expansion.

Zwicky and his colleague Walter Baade speculated that the explosive energy comes from gravity. Their idea was that

The other class of supernova involves the implosion of a star at least eight times as massive as the sun. This class is designated type Ib, Ic or II, depending on its observed characteristics.

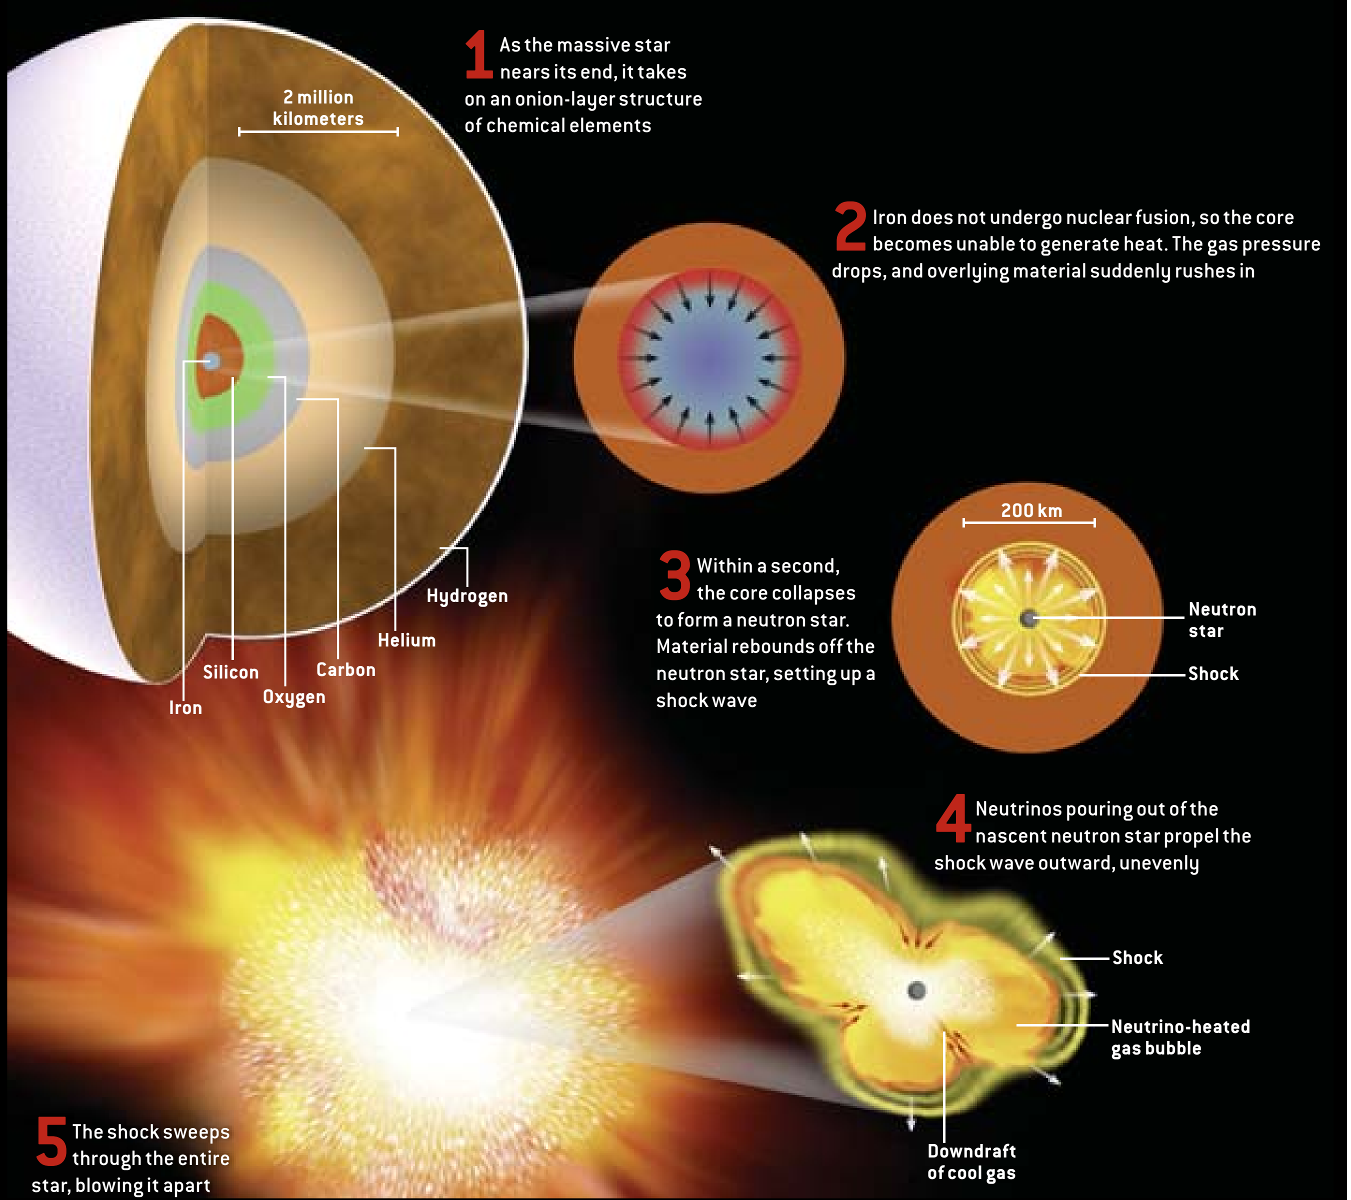

Accretion shock

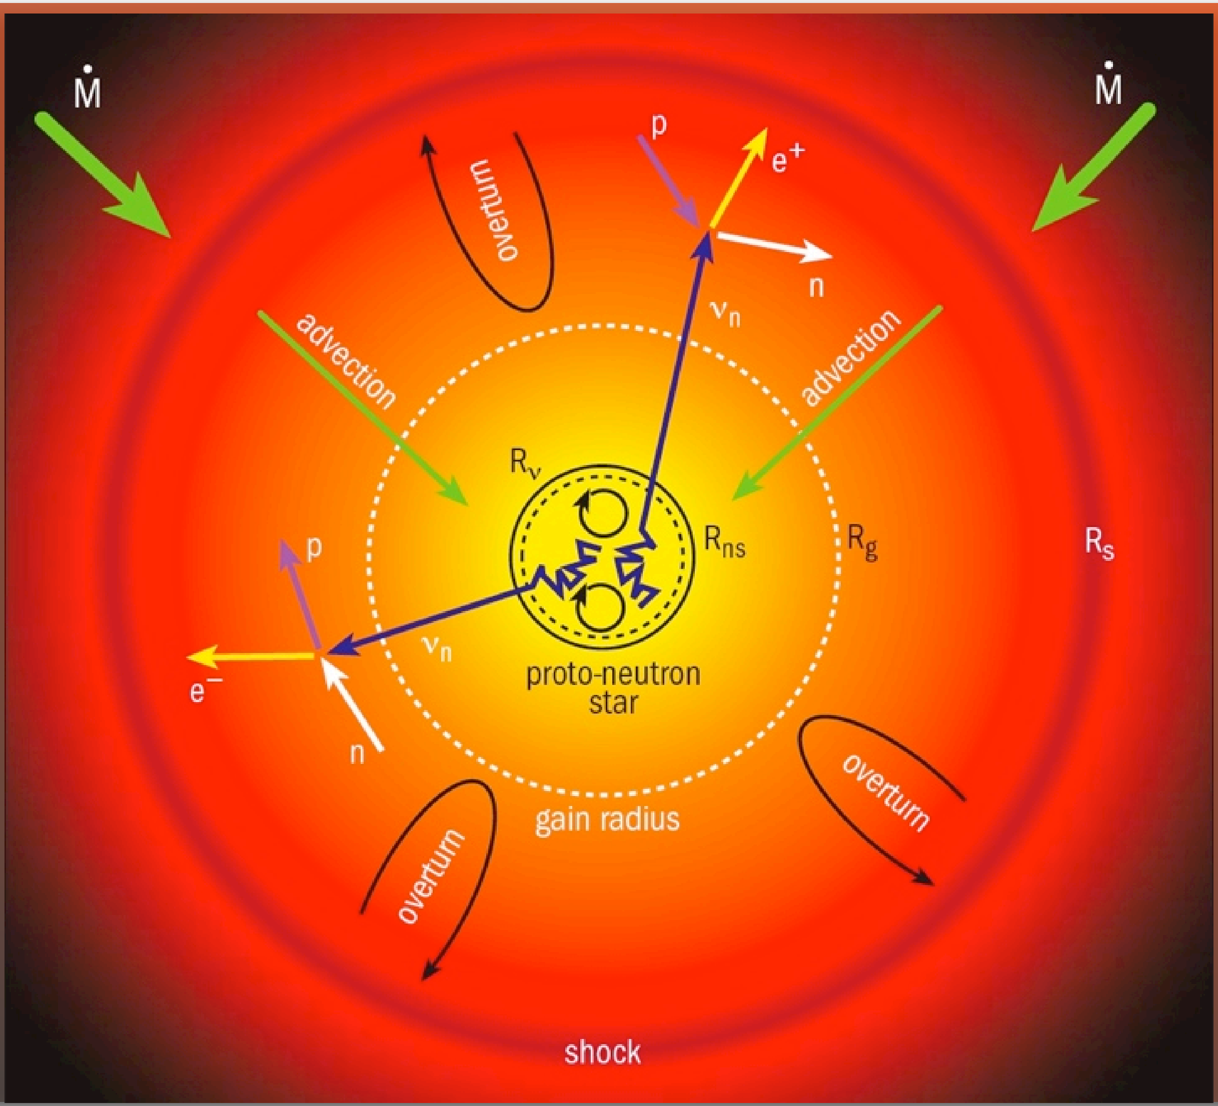

Current Workhorse

mCHIMERA

Ray-by-ray MGFLD transport (E_ν)
3D (magneto)hydrodynamics
150 species nuclear network

Possible Future Workhorse

bCHIMERA

Ray-by-ray Boltzmann transport (E_ν,θ)
3D (magneto)hydrodynamics
150-300 species nuclear network

The "Ultimate Goal"

GenASiS

Full 3D Boltzmann transport (E_ν,θ,ϕ)
3D (magneto)hydrodynamics
150-300 species nuclear network

- mCHIMERA is, well... a “chimera” of 3 separate, mature codes:
  - VH1 (MVH3)

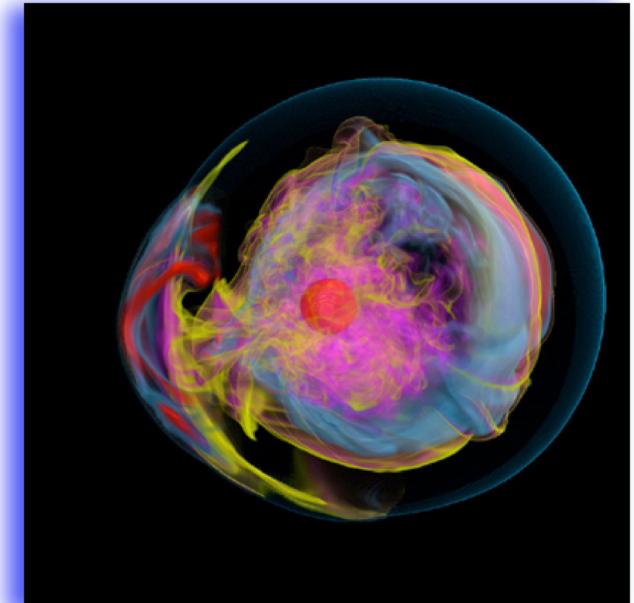

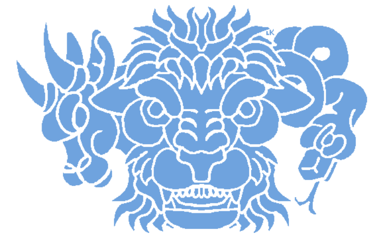

  - Multidimensional hydrodynamics
  - http://wonka.physics.ncsu.edu/pub/VH-1/
  - N. B. The CHIMERA version of VH1 is vastly different from the public version
    - non-polytropic EOS
    - 3D domain decomposition
    - **other sausage-like changes**

• MGFLD-TRANS

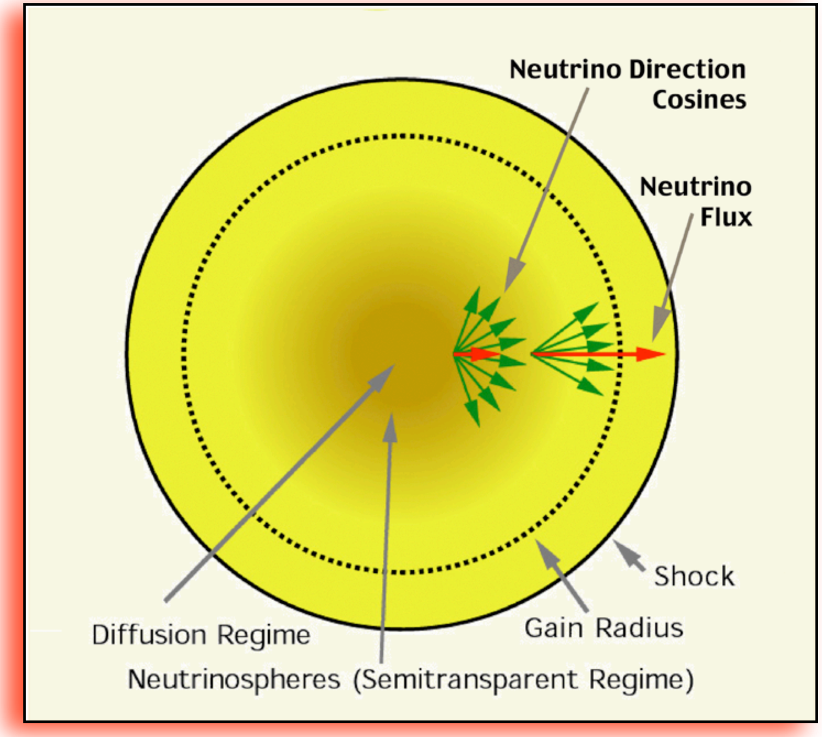

- Multi-group (energy) neutrino radiation hydro solver
    - GR corrections
    - 4 neutrino flavors with many modern interactions included
- flux limiter is “tuned” from Boltzmann transport simulations

- XNET

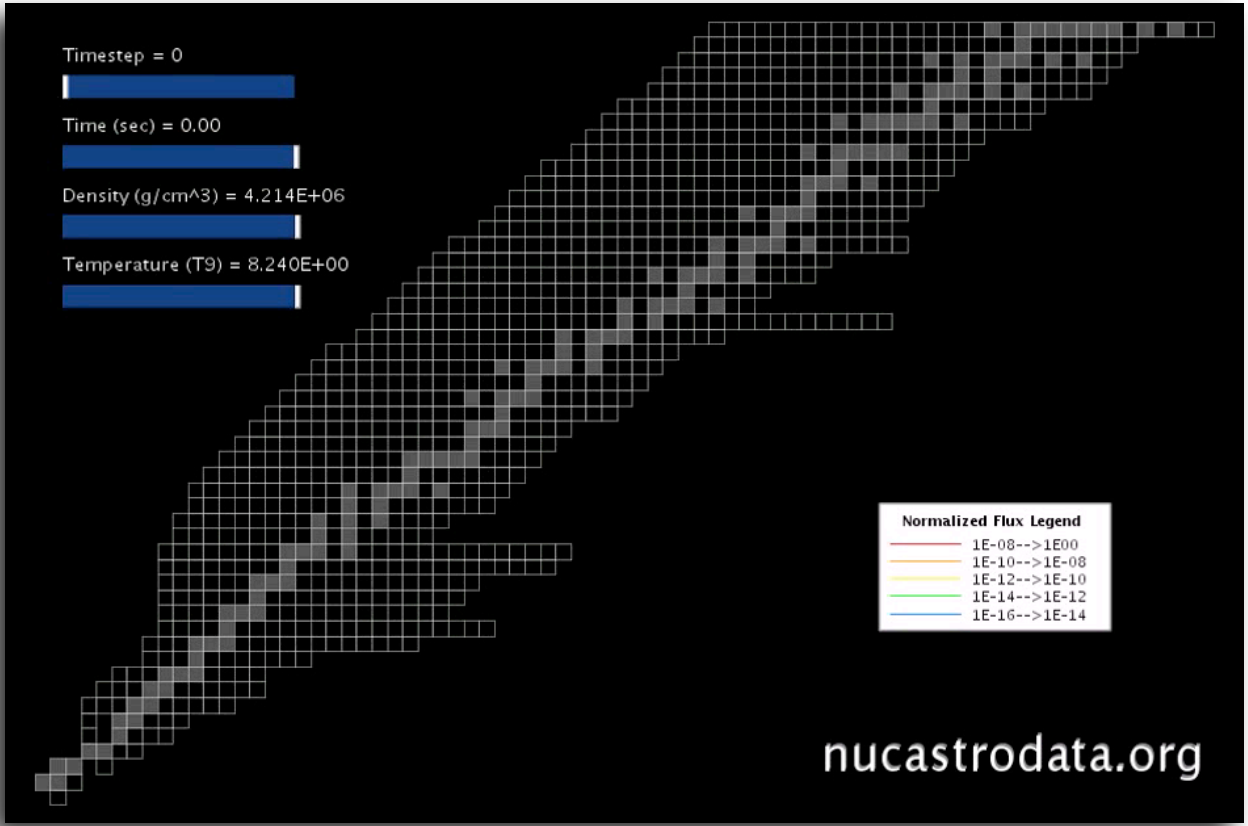

- Nuclear kinetics solver
    - Currently have implemented only an α network
    - 150 species to be included in future simulations
- Custom interface routine written for CHIMERA
- All else is 'stock'

How does CHIMERA work?

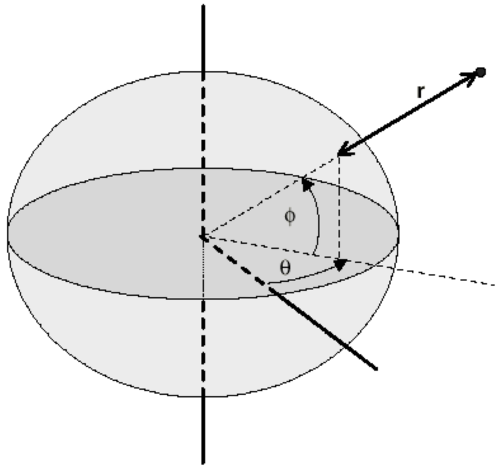

Big chunks of operator-split physics provide an opportunity to exploit the 250TF and the 1PF machine efficiently.

In fact, in a lot of different ways...

VH1/
MVH3

MGFLD-TRANS

XNET

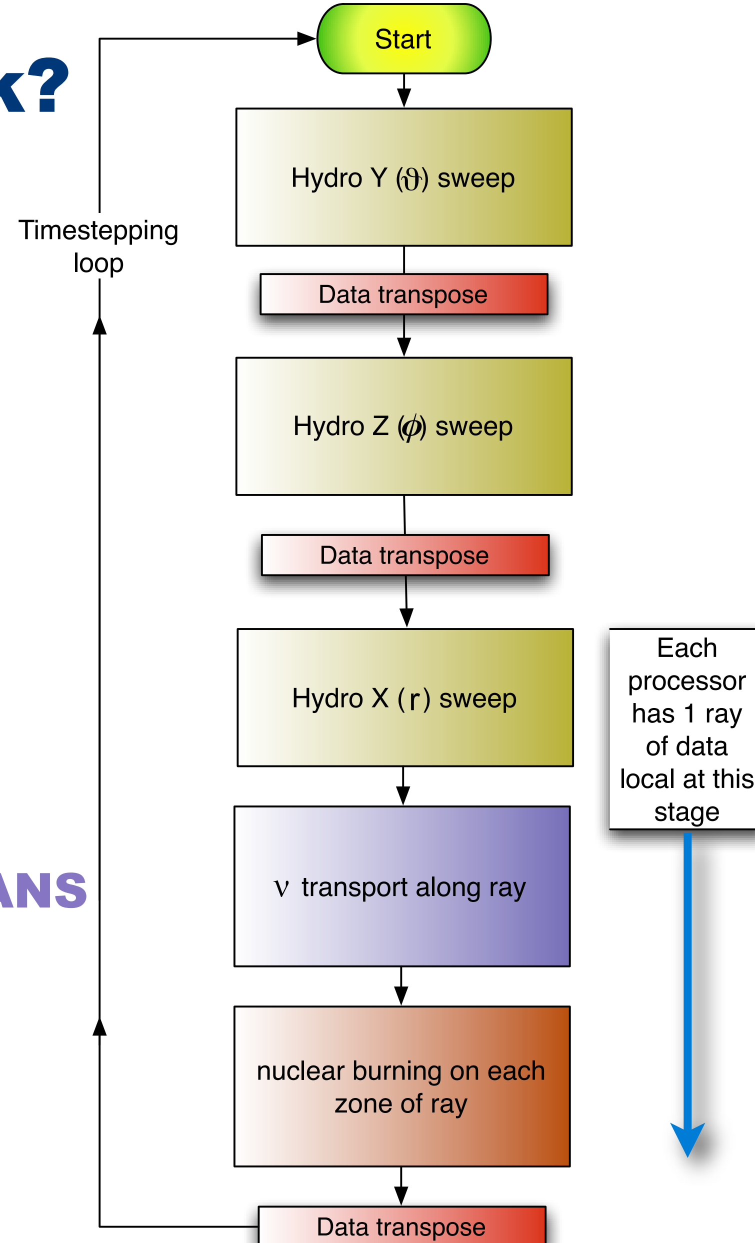

Hydrodynamics Scaling

- The hydrodynamics is a negligible component of the overall runtime (e.g. 0.04 s out of ~45 s total at 10K cores)

- On the 250TF, e.g., we'll be running this set of MPI tasks on ~20,000 sockets

- The "payload" for the collectives changes in CHIMERA

- How about Co-Array Fortran?

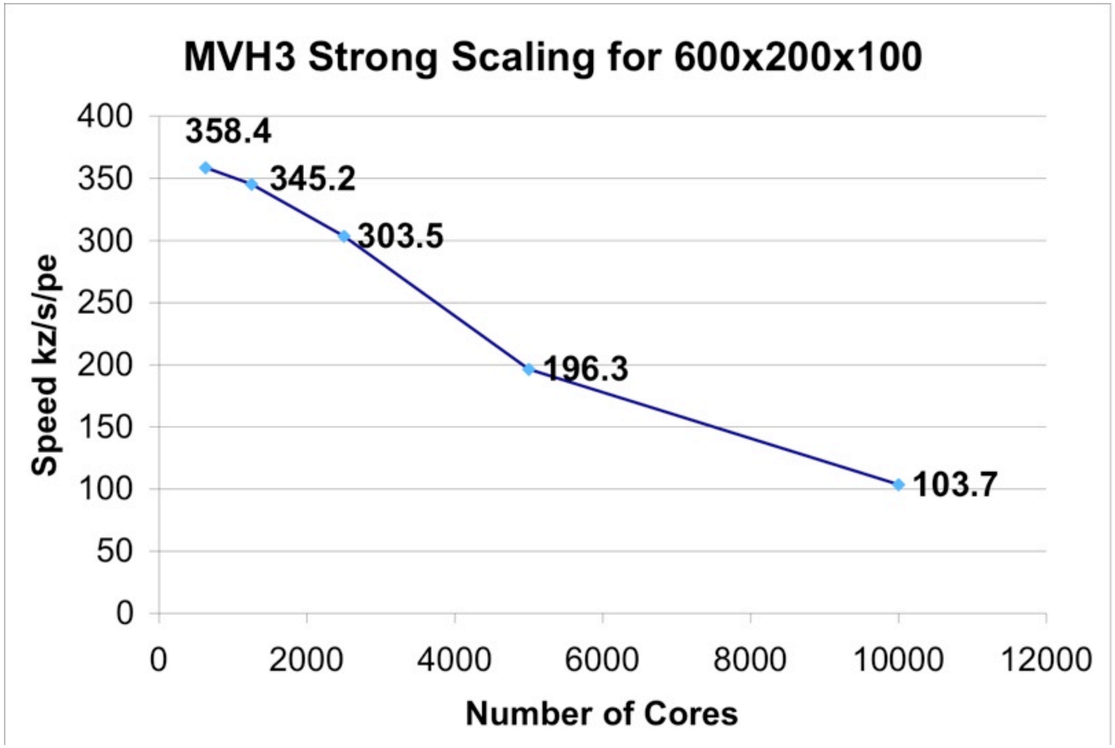

XNET performance and implementation

- XNET runs at ~50% of peak on a single XT4 processor
  - **Roughly 50% Jacobian build / 50% dense solve**
- 1 XNET solve is required per SPATIAL ZONE (i.e. hundreds per ray)
- Best load balancing on a node with OpenMP or a subcommunicator is interleaved

Transport module

- Transport solution is primarily 1 large-ish sparse solve
- Technology/Implementation for bCHIMERA:
  - **Boltzmann solver already uses threaded "ADI-like" preconditioner (D'Azevedo et al. 2005)**
    - **plan to move this preconditioner to mCHIMERA**
  - **Boltzmann transport solver also already parallelized along a ray, so subcommunicator across a "very-multi"-core socket is viable with little work (OR OpenMP OR Co-Array Fortran)**

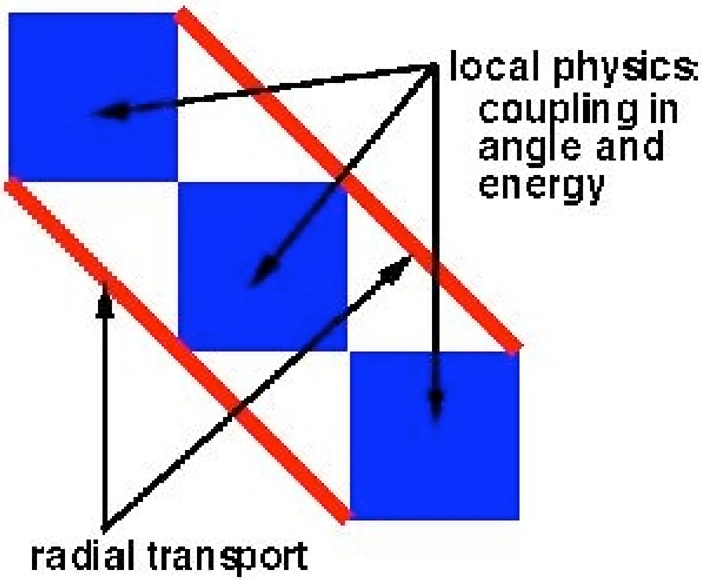

- Diagonal blocks are unsymmetric, ~80x80 for MGFLD, 500x500 for Boltzmann
- Overall system ~8000x8000, ~1% non-zeroes (MGFLD)

Summary

- The multi-scale and multi-physics characteristics of core-collapse supernova simulation makes it an ideal candidate for petascale (and beyond) computing
    - **huge scale contrasts**
    - **massive amounts of physics to be modeled**
    - **requires modern, sophisticated software infrastructure to make real progress**
- CHIMERA architecture allows realistic supernova simulations to be run on modern and near-future platforms
    - **fine-grain parallelism can be exposed in the neutrino transport and the nuclear burning**
    - **OpenMP? MPI sub-communicators? PGAS?**
    - **transport module can be swapped out**
    - **nuclear kinetics module can be made more sophisticated (e.g. QSE)**